

AMENDMENTS TO THE DRAWINGS

The attached "Replacement Sheets" of drawings include changes to Figures 1, 2, and 3a. The attached "Replacement Sheets," which include Figures 1, 2, and 3a, replace the original sheets including Figures 1, 2, and 3a.

Attachment: Replacement Sheets

REMARKS

Claims 1-10 and 15-16 are currently pending in the application. Independent claim 1 has been amended to distinguish the claimed invention over the cited references. Claims 2-7 and 10 have been amended to correct typographical errors, and are thus not limiting amendments. Claims 11-14 have been cancelled and claims 15-16 have been added. The Examiner is respectfully requested to reconsider and withdraw the rejections in view of the amendments and remarks contained herein.

SPECIFICATION

Applicants have amended the specification to fix typographical errors, include reference numerals, and include paragraph numbers. No new matter has been added. For the Examiner's convenience, a substitute specification is attached.

DRAWINGS

Figures 1, 2, and 3a were simply modified to fix typographical errors and to include missing reference numerals.

REJECTION UNDER 35 U.S.C. § 103

Claims 1, 11, and 12 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Park et al. (U.S. Pat. No. 5,872,771, hereinafter "Park") in view of Keshav ("An Engineering Approach to Computer Networking," 1997, hereinafter "Keshav"). Claims 11 and 12 have been cancelled, rendering their rejections moot, while the rejection with respect to claim 1 is respectfully traversed.

Park does not show a "first effective envelope associated with arriving traffic," a "second effective envelope associated with admitted traffic," or a "service curve by measuring departing traffic," as claim 1 recites. The scalar values that Park uses in admission control, such as cell loss rate, are not the same as the effective envelopes and service curve of the present invention, if only because the scalar values are not "functions of a time interval variable," as claim 1 recites.

Keshav teaches that "integrated service networks" are "likely to cost less than networks that provide a single quality of service." The present invention is applicable not only to integrated service networks, but to differentiated service networks and Multi-Protocol Label Switching. A teaching in favor of integrated service networks is irrelevant to an approach that applies equally to other types of network. Additionally, nothing Applicants are aware of in Keshav suggests using effective envelopes or service curves in admission control.

Applicants believe that new claims 14 and 15 are in condition for allowance, as they depend from base claim 1, which Applicants believe to be allowable.

Claims 2-8 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Park in view of Keshav in further view of Liebeherr et al. ("Effective Envelopes: Statistical Bounds on Multiplexed Traffic in Packet Networks," hereinafter "Liebeherr"). This rejection is respectfully traversed.

Liebeherr does not correct the deficiencies of Park as they relate to base claim 1. Accordingly, Applicants respectfully request that the Examiner withdraw the rejection.

Claims 9 and 10 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Park in view of Keshav in further view of Cruz et al. ("Scheduling for Quality of Service Guarantees via Service Curves," hereinafter -"Cruz"). This rejection is respectfully traversed.

Cruz does not correct the deficiencies of Park as they relate to base claim 1. Accordingly, Applicants respectfully request that the Examiner withdraw the rejection.

Claim 13 stands rejected under 35 U.S.C. § 103(a) as being unpatentable over Park in view of Keshav in further view of Mo et al. (U.S. Pat. No. 6,693,909). Claim 14 stands rejected under 35 U.S.C. § 103(a) as being unpatentable over Park in view of Keshav in further view of Taylor (U.S. Pat. No. 5,664,170). As Applicants have cancelled claims 13 and 14, these rejections are now moot.

CONCLUSION

It is believed that all of the stated grounds of rejection have been properly traversed, accommodated, or rendered moot. Applicant therefore respectfully requests that the Examiner reconsider and withdraw all presently outstanding rejections. It is believed that a full and complete response has been made to the outstanding Office Action, and as such, the present application is in condition for allowance. Thus, prompt and favorable consideration of this amendment is respectfully requested.

If the Examiner believes that personal communication will expedite prosecution of this application, the Examiner is invited to telephone the undersigned at (248) 641-1600.

Respectfully submitted,

Dated: 2 September 2005

By:



Gregory A. Stobbs, Reg. No. 28,764

HARNESS, DICKEY & PIERCE, P.L.C.
P.O. Box 828
Bloomfield Hills, Michigan 48303
(248) 641-1600

GAS/MRN

**MEASUREMENT-BASED ADMISSION CONTROL
UTILIZING EFFECTIVE ENVELOPES AND SERVICE CURVES**

BACKGROUND AND SUMMARY OF THE INVENTION

[0001] The present invention relates generally to network-based communication architecture. More particularly, the invention relates to providing controlled quality of service in packet-based networks through admission control.

[0002] Currently much of the Internet is designed to provide "best effort" service. The Internet Protocol (IP) is designed to deliver packets of information to an ultimate destination through any available internal links that enable delivery. The actual time it takes to deliver these packets depends on the route taken and the traffic congestion encountered en route. The original design of the Internet Protocol focused on providing ultimate delivery, with the actual time to achieve delivery being only a secondary consideration.

Deleted: in

[0003] As the uses for the Internet have grown, and as Internet traffic has expanded geometrically, the original emphasis on delivery over timing is being challenged. With multicast applications and streaming audio and video applications growing in popularity, packet delivery time has become a central focus. Quality of Service (QoS) is a term often used to describe the degree to which a communications network provides reliability, availability, and timely throughput. Quality of service addresses issues such as transmission speed, timeliness of packet delivery, amount of jitter a network introduces into packet streams, and the probability of outright packet loss. As businesses begin to rely more and more on their Internet presence, some have expressed willingness to pay more for higher

quality of service because the higher quality of service translates directly into a smoother, more responsive experience for their customers. Some Internet Service Providers thus offer different service level agreements through which they commit to provide different levels of service quality at different fee rates.

Deleted: service

Deleted: providers

[0004] There are many proposals for improving quality of service in packet-switched networks such as the Internet. Quality of service may be improved at the individual router level by making the routers faster and more intelligent. However, this also increases system cost. Other proposals address the quality of service issue at the network model level. At the network model level, the performance of individual routers are largely ignored; focus instead shifts to the aggregate performance of all routers in the network. One popular approach is to consider the aggregate network only in terms of its outer boundary or end-to-end performance. Using such an analytical approach, the performance of the entire network, and the quality of service it provides, can be largely controlled by the behavior of the routers occupying the edge of the network (that is, the routers at the ingress and egress points).

Deleted: the

[0005] There are several end-to-end network models for controlling quality of service (QoS) in popular use today. Among the leading models are *Integrated Services (IntServ)*, *Multi-Protocol Label Switching (MPLS)*, and *Differentiated Services (DiffServ)*. IntServ supports a per flow quality of service guarantee. It employs a relatively complex architecture in which resources are reserved by means of a signaling protocol that sets up paths and reserves resources. MPLS provides another quality of service guarantee approach that focuses on packet

Deleted: and

Formatted: Font: Italic

Formatted: Font: Not Italic

Formatted: Font: Italic

forwarding. Packets are assigned labels at the ingress of an MPLS-compatible domain. Subsequent classification, forwarding, and services for the packets are then based on those labels.

[0006] Models[I,II] such as IntServ and MPLS[II,III] address QoS on a *per connection* basis and can present scalability difficulties. To provide better scalability, other models have been proposed that address QoS on a *traffic aggregate* basis. DiffServ is one example of such a model. DiffServ provides quality of service guarantees to packet aggregates by marking packets differently to create different packet classes/aggregates that are entitled to different quality of service handling.

[0007] For the most part, aggregate traffic-based network models share a number of common concepts. They begin from a premise that the network can be characterized as having edge and core routers. The edge routers accept customer traffic (i.e., packets from any source outside the network). The core routers provide transit packet forwarding services among other core routers and/or edge routers. The edge routers control ingress of traffic and thus perform an important *admission control* function, by permitting or declining requests by outside traffic to enter the network. With the ultimate traffic flow being controlled by the edge routers[III] through admission control, the core routers simply need to differentiate traffic insofar as necessary to cope with transient congestion within the network itself. The network models may employ statistical multiplexing to maximize utilization of the core router resources.

Deleted: Core

[0008] Predicting and controlling traffic flow through a network at an aggregate level is a very complex problem. Admission control algorithms that overly restrict ingress waste internal core router resources. On the other hand, admission control algorithms that are too lax can flood the network with too much traffic, resulting in severe drops in quality of service. Numerous admission control algorithms have been proposed and while some have been quite ingenious, there remains a great deal of room for improvement.

Deleted: the

Deleted: drop

Deleted: There have been n

Deleted: . While

Deleted: of these

[0009] For example, Centinkaya and Knightly, in an article entitled "Egress Admission Control," IEEE Info. Comm. 2000, describe a framework for scalable quality of service provisioning in which admission control decisions are made at the egress routers based solely on aggregate measurements obtained at the egress router. They introduce a measurement-based service envelope as a way to adaptively describe the end-to-end service received by a traffic class.

Deleted: scalable

[0010] For a more complete understanding of the invention, its objects and advantages, refer to the following description and to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] Figure 1 is a network diagram illustrating an exemplary end-to-end network model, useful in understanding the admission control algorithm of the present invention;

[0012] Figure 2 is a block diagram illustrating the pertinent parameters associated with arriving traffic and admitted traffic in a network;

Deleted: [0010] The present invention is an improvement upon the above technique in which a measurement-based admission control algorithm is based on the global effective envelopes of an arriving traffic aggregate and upon the service curves of the corresponding departing traffic aggregate. The admission control algorithm is well adapted to a variety of different network models. Although the invention will be principally described in the context of a multi-router network, the concepts employed are equally applicable to autonomous network systems and even switching devices. In a switching device, for example, the switch hardware/software provides several input ports and output ports. The present invention may be used to control flow of traffic among these input and output ports. Therefore, the invention is applicable to "networks" of different architectural granularities, ranging from single device switches to large, multiple device computer network domain.

[0013] Figure 3a is a flowchart diagram of the presently preferred admission control algorithm, illustrating global effective envelopes;

[0014] Figure 3b is a similar flowchart diagram of the presently preferred admission control algorithm, illustrating local effective envelopes;

[0015] Figure 4 is a graph illustrating the calculation of delay bound;

[0016] Figure 5 is a timing diagram illustrating an example of a measured time interval, where $T = 3\tau$ and $M = 2$;

[0017] Figure 6a and 6b are graphs illustrating measured arriving traffic for time interval ≤ 100 ms;

Deleted:

Deleted: kr

[0018] Figures 7a and 7b are graphs illustrating measured service envelope for time interval ≤ 100 ms; and

[0019] Figures 8a, 8b, and 8c are graphs of average link utilization as a function of delay bound for three different scenarios, for delay bound $d \leq 100$ ms.

DESCRIPTION OF THE PREFERRED EMBODIMENT

[0020] A measurement-based admission control algorithm is based on the global effective envelopes of an arriving traffic aggregate and upon the service curves of the corresponding departing traffic aggregate. Such an admission control algorithm is well adapted to a variety of different network models. Although the algorithm will be principally described in the context of a multi-router network, the concepts employed are equally applicable to autonomous network systems and even switching devices. In a switching device, for example, the switch hardware/software provides for several input ports and output ports. The present

invention may be used to control traffic flow among these input and output ports. Therefore, this algorithm is applicable to "networks" of different architectural granularities, ranging from single-device switches to large, multiple-device computer network domains.

[0021] Referring to Figure [1], an exemplary network is illustrated at 10. Network 10 includes a plurality of edge routers, such as router 12 and router 14, that are in turn connected to a plurality of core routers 16. For purposes of illustration, edge router 12 serves as the ingress node 18, and edge router 14 serves as the egress node 20. Arriving traffic enters through the ingress node 18 and departing traffic leaves through the egress node 20, as illustrated.

[0022] For purposes of illustration, arriving aggregate traffic 22 comes from a second network 24 via edge router 26. For purposes of illustration, it will be assumed that a traffic conditioner or regulator function 28 [1] limits the flow of arriving aggregate traffic 22 to within predefined bounds. The traffic conditioner 28 may be located within the second network 24 (such as in edge router 26). The traffic conditioner or regulator 28 may be implemented at the ingress node 18 to force arriving traffic to comply with the traffic specification submitted by the sender. The traffic conditioner or regulator 28 may be modeled as a leaky bucket that provides a simple metering function. Alternatively, the traffic conditioner 28 may provide more sophisticated traffic shaping functions where the temporal characteristics of a traffic class are selectively modified by delaying local packet forwarding. For purposes of illustration, departing aggregate traffic 30 exits edge router 14 and enters network 32.

Formatted: Font: Not Bold

Deleted: the

Formatted: Font: Not Bold

Formatted: Font: Bold

Deleted: the

Deleted: the edge router 26 includes

Deleted: ing

Formatted: Font: Not Bold

Formatted: Font: Bold

Formatted: Font: Not Bold

Formatted: Font: Not Bold

Deleted: the

[0023] Network 10 implements an admission control function 34 to determine which of the packets of arriving aggregate traffic 22 are admitted to the network 10 and which are blocked. The admission control function 34 may be localized in a single router, such as edge router 12. Alternatively, the admission control function can be distributed across the network, with the function being performed by multiple routers within the network 10. The admission control function 34 implements an admission control algorithm that is based upon a global effective envelope 36 of the arriving aggregate traffic 22 and upon service curves 38 of the departing aggregate traffic 30.

Deleted: plural

Deleted: In accordance with the invention

Deleted: the

Deleted: the

Deleted:

[0024] Figure 2 illustrates network 10 with arriving traffic 22, admitted traffic 46, and departing traffic 30 being shown in greater detail. The figure introduces terms that will be useful in understanding the algorithm.

Deleted: the

[0025] Associated with arriving traffic 22 is a deterministic envelope 42. The deterministic envelope may be characterized in terms of three variables: peak traffic rate, average traffic rate, and burst size parameter. The arriving traffic 22 also is assumed to have an associated quality of service requirement 44. Typically each class of service would have its own quality of service requirement. Thus multiple QoS requirements are illustrated at 44 in Figure 2. The QoS requirements may be characterized by the following variables: probability of delay violation and delay bound.

Deleted: some

Deleted: of the

Deleted: presently preferred

Deleted: The arriving traffic 22 is assumed to be supplied through a traffic conditioner or regulator 28, which may be implemented at the ingress node 18 to enforce the arriving traffic to comply with the traffic specification submitted by the sender. The traffic conditioner or regulator may be modeled as a leaky bucket that provides a simple metering function. Alternatively, the traffic conditioner may provide more sophisticated traffic shaping functions where the temporal characteristics of a traffic class are selectively modified by delaying local packet forwarding.

[0026] The admitted traffic 46 may be differentiated from the arriving traffic 22, in that the admitted traffic 46 has already passed the admission control point and is thus currently flowing in network 10. When the admitted traffic 46 leaves

network 10, it is called departing traffic 30. The departing traffic 30 has associated with it a service curve or service envelope 48 that is measured by monitoring each packet's arrival time and departure time.

Deleted: represents

[0027] One variation of the admission control algorithm uses the associated deterministic envelope 42 of the arriving traffic 22 and the associated service envelope 48 of the departing traffic 30 in making a decision whether to admit or deny arriving traffic 22. Figures 3a and 3b show the presently preferred algorithm in flowchart form. A more rigorous mathematical derivation of the preferred algorithm is presented below. The flowcharts of Figures 3a and 3b make reference to specific equations from the detailed mathematical derivation, where applicable. A discussion of the flowchart will be presented first, as a way of introducing the detailed mathematical derivation that follows. Specifically, Figure 3a applies to global effective envelopes; Figure 3b applies to local effective envelopes. The content of these two figures is similar; hence a detailed discussion of Figure 3a is presented below and may be considered applicable to Figure 3b, except where the mathematical derivation differs, as will be evident by comparing the Figures.

Deleted: The

Deleted: of the presently preferred embodiment

Deleted: to meet the desired QoS requirements

[0028] Referring to Figure 3a, the preferred algorithm begins at step 60 with a request to be admitted to the network. As part of the request, the incoming traffic submits its traffic characterization, which is expressed in terms of characterization parameters 62. The specific parameters used by the presently preferred embodiment are:

- Deterministic Envelope $A_j(\tau)$;
- Quality of Service Requirements, QoS;

- Delay Bound, d ;
- Probability of Delay Violation, ϵ .

[0029] Next, at step 64, the algorithm constructs a global effective envelope for the arriving traffic, based on the submitted parameters 62. In the mathematical derivation that follows, the global effective envelope for arriving traffic is characterized as G_{new} . The G_{new} envelope is depicted in Figure 3 at 66. As indicated by the dashed lines and bracketed notation, step 64 is performed using equations 3.8-3.11.

[0030] The algorithm continues at step 68, where a global effective envelope for the admitted traffic[1,1] is constructed based on measured statistics. Specifically, the measured statistics, listed at 70, include the average and the variance of the amount of arriving aggregate traffic. In the detailed mathematical derivation which follows, the global effective envelope for admitted traffic is designated as $G_{N_s}^g$. In Figure 3, this $G_{N_s}^g$ envelope is shown at 72.

Deleted: Next the

Deleted: to construct

[0031] The final admission control decision, step 74, is made based on an analysis of both global effective envelopes 66 and 72, along with the service curve 40 calculated using equation 4.5 and the algorithm shown in Table I below. Specifically, the admission control decision will admit incoming traffic if:

$$G_{new} + G_{N_s}^g \leq S^g$$

For a more precise statement of the above admission control condition, see the following detailed description and, in particular, Egn. (3.12).

Deleted: detailed

Detailed Description of Preferred Admission Control Algorithm

[0032] A detailed description of the preferred algorithm will now be provided below. For purposes of illustration, a Differentiated Service (DS) or Diffserv architecture will be assumed. The invention is not limited to the DS architecture inasmuch as the admission control algorithm of the invention may be used with a variety of network architectures.

Deleted: DiffServ

[0033] In order to provide Quality of Service (QoS), the Integrated Services (IS) architecture uses a signaling protocol to set up resource reservations at all routers along the path. Since this approach requires that all routers have to keep per-flow state and process the per-flow resource reservations, this approach has known scalability problems. Differentiated Service (DS) is another approach to provide QoS, but without the restriction on scalability. Routers need only implement scheduling and buffering mechanisms and apply them based on the DS Code Point (DSCP) presented in the header of arriving packets. The end-to-end QoS service across several domains is built by combining the per domain behaviors (PDBs) of individual DS domains. PDBs, an active area of research, are constructed by an appropriate deployment of per-hop forwarding behaviors (PHBs). Several DS domains can exist within a single autonomous system (AS). DS addresses scalability by providing QoS guarantees on an aggregate traffic level.

Deleted: use

Deleted: to

Deleted: Those

Deleted: in their terms

Deleted: The definition of PDBs is an active area of research, with the first few drafts currently available). Hence

Deleted: aggregate

[0034] As discussed above, one of the critical requirements for service provisioning across a DS domain is [[a]] the presence of an effective Connection Admission Control (CAC). Such admission control can be utilized while

establishing dynamic service level agreements (SLAs) for premium service across DS network, or consulted upon re-negotiations of static SLAs. A Bandwidth Broker (BB) entity of a DS domain is a typical place for such admission control. The difficulty involved in designing an effective CAC[.] is that many solutions lead toward topology-dependent implementations, thus limiting their scope and increasing their complexity.

Deleted: the

[0035] To implement admission control for DS without scalability problems, several researchers have proposed various forms of Endpoint CAC. In this design, the endpoints of a DS domain (ingress/egress routers) perform admission control and resource management. The routers measure traffic in the network to detect the level of available resources and admit a new connection if, and only if, the detected level of resource is sufficiently high.

[0036] The preferred embodiment provides a measurement-based CAC algorithm, which provides statistical QoS guarantees on a traffic aggregate level in a DS domain. The ingress traffic arriving at the boundary of a DS domain is assumed to be conditioned by the upstream domains, policed by our domain to adhere to the rules of an established SLA, and further marked with the appropriate DSCP. Leaky-bucket based policers, shapers and markers are widely proposed. As the DS architecture extends some PDB for a traffic aggregate, we call this traffic aggregate a class in a subsequent discussion (different from DS class). We target the PDBs, which offer delay guarantees to a classified traffic aggregate given both delay and delay violation probability. We describe the deterministic envelopes, i.e., arriving traffic envelopes and service

Deleted: described

envelopes, and the bound on performance characteristics that can be obtained~~[,]~~ from those envelopes. We then derive a measurement-based CAC to establish dynamic SLA and provide statistical guarantees for a variety of service classes by exploiting the statistical traffic envelopes, called *global effective envelopes* and *service curves*. To obtain these envelopes, we use the measurement approach. The measurement of the service curves allows our admission control condition to work without the knowledge of the topology of the DS domain and cross traffic in the domain.

Deleted: which

[0037] Let us consider a traffic arriving to a network domain. The traffic enters the domain at an ingress node 18 and destined to an egress node 20 (see Figure 1). In the following we consider a continuous-time fluid flow traffic model. An arriving traffic from different flows is differentiated by its QoS requirements. Here we consider the probability of delay violation and the delay bound as the QoS requirements of interest. An arriving traffic with the same QoS requirements, i.e., the same delay bound and probability of delay violation requirements, is treated as a service class. Hence, traffic arriving to the domain is partitioned into Q classes. We use N_q to denote the number of flows within a class q .

Deleted: low

Deleted: traffic in

[0038] The arriving traffic from the flow j of the class q in an interval (t_1, t_2) is denoted as $A_j^q(t_1, t_2)$. We assume that traffic arrival from each flow j is regulated by a traffic conditioner and upper bounded by a deterministic subadditive envelope $A_j^{q*}(\tau)$ as follows:

$$A_j^q(t, t + \tau) \leq A_j^{q*}(\tau) \quad \forall t \geq 0, \forall \tau \geq 0. \quad (2.1)$$

[0039] The arriving traffic is characterized by the deterministic envelope with a set of parameters. The most commonly used traffic regulators are leaky buckets with a peak rate enforcer. The traffic on flow j is characterized then by three parameters (P_j, σ_j, ρ_j) with a deterministic envelope given by

Deleted: is

$$A_j^q(\tau) = \min\{P_j^q \tau, \sigma_j^q + \rho_j^q \tau\} \quad \forall \tau \geq 0 \quad (2.2)$$

where $P_j^q \geq \rho_j^q$ is the peak traffic rate, ρ_j^q is the average traffic rate, and σ_j^q is the burst size parameter.

[0040] As a generalization of the peak-rate enforced leaky bucket, the traffic on flow j may be characterized by a set of parameters $\{\sigma_{jk}^q, \rho_{jk}^q\}_{k=1, \dots, K_j}$, with a deterministic envelope

$$A_j^q(\tau) = \min_{k=1, \dots, K_j} \{\sigma_{jk}^q + \rho_{jk}^q \tau\} \quad (2.3)$$

where K_j^q is the number of leaky-bucket pairs.

[0041] It is also noted that flows from the same service class can be characterized by different deterministic envelopes, but still have the same QoS requirements.

Deleted: can

[0042] We use $A_j^{qout}(t_1, t_2)$ to denote the traffic from flow j of the class q leaving the domain at the egress node in an interval (t_1, t_2) . We then define the aggregate traffic of class q as follows: the aggregate of arriving traffic $A_{N_q}(t_1, t_2) = \sum_{j \in N_q} A_j^q(t_1, t_2)$ and the aggregate of departing traffic $A_{N_q}^{out}(t_1, t_2) = \sum_{j \in N_q} A_j^{qout}(t_1, t_2)$. For a time interval of length β , we define the empirical envelope, $E_{N_q}(\tau, \beta)$ of an arriving aggregate N_q of flows from class q , which is the

upper bound of the arriving aggregate in any time interval of length $\tau \leq \beta$ as follows:

$$E_{N_q}(\tau, \beta) = \sup_{(t, t+\tau) \subset \beta} A_{N_q}(t, t+\tau) \quad (2.4)$$

[0043] In the following, we will focus our discussion on the aggregate traffic of a single service class and drop the index 'q'. We will introduce the definition of a service curve and its application. We assume that there is no traffic in the domain at time $t = 0$. Therefore, for an aggregate traffic, the amount of backlogged traffic in the domain at any time t , $B(t)$, is given by

$$B(t) = A_N(0, t) - A_N^{out}(0, t) \geq 0. \quad (2.5)$$

Also the ingress-to-egress delay suffered by the aggregate traffic at time t is denoted by $d(t)$ and is defined as follows:

$$d(t) = \min\{z : z \geq 0 \text{ and } A_N(0, t) \leq A_N^{out}(0, t+z)\} \quad (2.6)$$

Definition of Service Curve

[0044] Let $S(\cdot)$ be a non-decreasing function, with $S(0) = 0$. We say that the domain guarantees the service curve $S(\cdot)$ for the aggregate traffic if for any time t , there exists $s \leq t$ such that $B(s) = 0$ and $A_N^{out}(s, t) \geq S(t-s)$.

[0045] It can be seen that $S(t-s)$ specifies the lower bound of the amount of the aggregate traffic needed to be served and depart from the domain during some interval (s, t) , where the domain is empty at time s .

[0046] From the definition of service curve, the upper bound of delay can be obtained as follows. Consider a busy period of length β . Let s denote the starting

time of the busy period, that is, $B(s) = 0$ and $A_N(0, s) = A_N^{out}(0, s)$. We assume that the domain guarantees the service curve $S(\cdot)$, that is $A_N^{out}(s, t) \geq S(t - s)$, where $s \leq t + d \leq \beta$. From the definition of service curve, we have $A_N^{out}(s, t + d) = A_N^{out}(0, t + d) - A_N(0, s) \geq S(t + d - s)$. From Eqn. (2.6), thus, $A_N(0, t) \geq A_N^{out}(0, t + d) \geq A_N(0, s) + S(t + d - s)$. Hence, from Eqn. (2.4),

$$A_N(s, t) \leq E_N(t - s; \beta) \geq S(t + d - s) \quad (2.7)$$

[0047] Defining $\tau = t - s$, it follows that, for all time $t \geq 0$, the delay defined in Eqn. (2.6) is upper bounded as follows:

$$\begin{aligned} d(t) &\leq \min\{z : z \geq 0 \text{ and } E_N(\tau, \beta) \leq S(\tau + z)\} \\ &\leq \max_{0 \leq \tau \leq \beta} \min\{z : z \geq 0 \text{ and } E_N(\tau, \beta) \leq S(\tau + z)\} \end{aligned} \quad (2.8)$$

From Eqn. (2.8), it is noted that the upper bound on delay is the maximum horizontal distance between the arriving aggregate envelope ($E_N(\tau, \beta)$) and the service curve $S(\tau)$ (see Figure 4). It is also noted that β is upper bounded by β_{\max} , which is the maximum busy period of the aggregate traffic and can be determined as: $\beta_{\max} = \inf\{\tau \geq 0 \mid \sum_{j \in N} A_j^*(\tau) \leq S(\tau)\}$. Therefore, we will use $\beta = \beta_{\max}$ in the following discussions.

[0048] The preferred admission control algorithm exploits the deterministic envelopes and the service envelopes to obtain the upper bound of the delay experienced by an arriving traffic based on the worst-case scenario where it is assumed that all flows send traffic with their peak rate simultaneously.

[0049] In the next section, following the concept of the described envelopes, we will derive a statistical service where a small portion of traffic is allowed to

violate the delay QoS requirement. We will use so called global effective envelopes to characterize arriving aggregate traffic for each class. Using these envelopes allows us to derive a CAC for statistical services which can be created in the same fashion as for deterministic services.

Deleted: perform

Deleted:

Probabilistic Guarantees with Effective Envelopes and Service Curves

[0050] In this section, we derive the probability of delay violation based on the upper bound of the delay shown in the previous section. We then obtain an admission control condition for a single service class by applying the definition of local and global effective envelopes presented in R. Boorstyn, A. Burchard, J. Liebeherr, and C. Oottamakorn[1], Effective envelopes: Statistical bounds on multiplexed traffic in packet networks, *Proceedings of IEEE Info. Comm. 2000*, Tel Aviv, Israel, March 2000 (hereinafter "Boorstyn, et al."). We finally discuss the admission control conditions for multiple service classes.

Deleted: envelope

Deleted: .

Deleted: in

Deleted: INFO

Deleted: COMM

Probabilistic Guarantees with Local Effective Envelopes and Service Curves

[0051] Again we will first consider a single service class, allowing us to drop the index 'q'. An aggregate of arriving traffic from the same service class $A_N(t_1, t_2)$ has the same probability of delay violation, ε , and delay bound requirement, d . $A_N(t_1, t_2)$ are upper bounded by a family of nonnegative random processes $A_N(t_1, t_2)$, which have the following properties: (1) Stationarity: the statistical properties of $A_N(t_1, t_2)$ do not change with time, and (2) Independence: The $A_N(t_1, t_2)$ are stochastically independent among all service classes.

Deleted: local

Deleted: envelopes

Deleted:

Deleted: and we

Deleted: has

[0052] Let $D_N(\tau)$ denote the ingress-to-egress delay experienced by a bit of an aggregate traffic of the N flows arriving at the domain at time τ . We assume that $D_N(\tau)$ is a random variable. For a given delay bound d , the domain guarantees that the ingress-to-egress delay of any arriving bit will not exceed the delay bound d ; that is, for all time $t \geq 0, 0 \leq \tau \leq B_{\max}$, from Eqn. (2.8), $D_N(\tau) = \min\{z : z \geq 0 \text{ and } A_N(t, t+\tau) \leq S(t+z)\} \leq d$ or, for all time $t \geq 0, 0 \leq \tau \leq B_{\max}, A_N(t, t+\tau) \leq S(t+d)$. The delay violation occurs if $\exists D_N(\tau)$ such that

$$D_N(\tau) \geq d$$

or, $\exists \tau$ such that

$$A_N(t, t+\tau) > S(t+d)$$

[0053] With Eqn. (2.8), the probability that the arriving traffic experiences a delay violation is required to be bounded by ε and we have, for all time $t \geq 0, 0 \leq \tau \leq B_{\max}$,

$$\Pr[D_N(\tau) > d] = \Pr[A_N(t, t+\tau) > S(t+d)] \leq \varepsilon \quad (1)$$

When a set of new flows requests the same service guarantees from a DS domain, the CAC uses the admission control test to ensure that the requested service guarantees and other flows from all service classes can be simultaneously satisfied. From Eqn. (1), we can obtain the admission control condition for a single service class by using the definition of local effective envelope as follows:

Deleted: again

Deleted: which

Formatted: Justified

Formatted: Centered

Deleted: request

Deleted: The

Definition (Local Effective Envelope)

[0054] A local effective envelope for $A_N(t, t + \tau)$ is a function L_N that satisfies for all $\tau \geq 0$ and all t

$$\Pr[A_N(t, t + \tau) \leq L_N(\tau, \epsilon)] \geq 1 - \epsilon$$

It can be seen that a local effective envelope provides a probabilistic bound for the aggregate traffic $A_N(t, t + \tau)$ for any specific time interval of length τ .

Lemma

[0055] Given a set of flows that is partitioned into M classes, each class contains N_m flows with aggregate arrivals A_{N_m} . Let $L_{N_m}(\tau, \epsilon)$ be the local effective envelope for class m . Then the following inequality holds: given x ,

$$\text{If } \sum_m L_{N_m}(\tau, \epsilon) \leq x, \text{ then for all } t, \Pr\left[\sum_m A_{N_m}(t, t + \tau) > x\right] \leq M \cdot \epsilon.$$

We divide flows from the same service class into two categories, i.e. $M = 2$. The first group is the set of flows, N_1 that are already admitted to the domain. The second group is the new set of flows, K , requesting the QoS service from a DS domain. We denote L_N as the local effective envelope of the first group and L_{new} as the local effective envelope of the second group.

Admission control condition for a single service class

[0056] Using the definition of the local effective envelope and the lemma, from Eqn. (1) we have that the arriving traffic has a delay violation with probability $< \epsilon$ if

$$L_N(\tau, \varepsilon/2) + L_{new}(\tau, \varepsilon/2) \leq S(\tau + d), \text{ for } 0 \leq \tau \leq B_{max} \quad (2)$$

[0057] If, for the same service class, the set of new flows are regulated by different deterministic envelopes A_j^* , each set of flows sharing the same A_j^* have a different envelope $L_{new}^j(\cdot; \varepsilon_j)$, where $\sum_j \varepsilon_j = \varepsilon/2$. Therefore Eqn. (2) becomes

$$L_N(\tau, \varepsilon/2) + \sum_j L_{new}^j(\tau, \varepsilon_j) \leq S(\tau + d), \text{ for } 0 \leq \tau \leq B_{max} \quad (3)$$

[0058] The CAC in Eqn. (2) does not depend on the topology of the domain and the scheduling algorithms within the domain, but the measurement-based traffic-characterized envelopes, which can be measured at an ingress and an egress node in a DS domain.

Local Effective Envelope using Measured Moment

[0059] We consider a set of admitted flows, N . We characterize the distribution of the aggregate traffic using the Central Limit theorem, which does not require knowledge of the underlying distributions of each flow. An approximate local effective envelope of the aggregate traffic can be obtained from the Central Limit Theorem as follows [1]. From the definition of the local effective envelope, we set $\Pr[A_N(t, t + \tau) \geq x] \approx \varepsilon/2$ and we obtain,

$$L_N(\tau, \varepsilon/2) \approx \overline{X}(\tau) + z\sigma(\tau) \quad (4)$$

where z has the approximate value $z \approx \sqrt{\log(2/\varepsilon)}$, and $\overline{X}(\tau)$ and $\sigma^2(\tau)$ are measured average and variance of the amount of aggregate traffic, respectively.

We will discuss those measurements in the next section.

[0060] Since the number of new flows may be few, the Central Limit theorem may not be the appropriate choice for constructing their local effective

Deleted: [0057]

Deleted: Remark.

Formatted: Font: (Default) Arial

Deleted: for

Deleted: which have

Deleted: it is required to construct their

Deleted: [0058]

Deleted: the

Deleted: just a

Deleted: flows

envelope, i.e., the bound obtained may not be tight enough. Instead, we will use the Chernoff bound, which provides a more rigorous bound on the aggregate traffic. The local effective envelope of a set of new flows, K , can then be obtained by using the Chernoff bound, and is given as follows:

Recall the Chernoff bound for a random variable Y :

$$\Pr[Y \geq y] \leq e^{-y} E[e^y] \quad \forall y \geq 0 \quad (5)$$

We have

$$\Pr[A_K(s, \tau) \geq Kx] \leq e^{-Kx} M_K(s, \tau) \quad (6)$$

where $M_K(s, \tau)$ is the moment generating function of the aggregate traffic.

[0061] In Boorstyn, et al., $M_K(s, \tau)$ is derived and given in terms of the deterministic envelopes of the flows, which are the leaky bucket function in our case. s is chosen so that Eqn. (6) is minimal. These values of s and $M_K(s, \tau)$ yield

$$L_{new}(\tau, \epsilon/2) = K \min\{x, A_K^*(\tau)\} \quad (7)$$

$$\left(\frac{\rho\tau}{x}\right)^{\frac{x}{A_K^*(\tau)}} \left(\frac{A_K^*(\tau) - \rho\tau}{A_K^*(\tau) - x}\right)^{1 - \frac{x}{A_K^*(\tau)}} \leq (\epsilon/2)^{1/K} \quad (8)$$

where $A_K^*(\tau)$ is a deterministic envelope for each new flow.

Admission control condition for multiple service classes

[0062] The admission control condition in Eqn. (2) does not guarantee that QoS requirements of other service classes using the same ingress-to-egress path will be satisfied if a new set of flows is admitted into the domain. To be able to provide that guarantee, without loss of generality, we assume that the new flows have the highest priority, and we test the admission control condition for each of the service classes. This can be shown as follows. For each service class we obtain an

Deleted: [0061]
Deleted: term
Deleted:
Deleted: that
Deleted: is
Deleted: With
Deleted:
Deleted: it
Deleted: yield
Deleted: s

Deleted: provide the
Deleted: use
Deleted: their
Deleted: s
Deleted: all

arriving aggregate traffic envelope $L_{N_q}^q(\tau, \varepsilon_q/2)$ and a service envelope $S^q(\cdot)$.

From Eqn. (2), the new flows with an arriving envelope $L_{new}(\tau, \varepsilon_q/2)$ will be admitted if the following condition is satisfied: for all service classes q ,

$$L_{N_q}^q(\tau, \varepsilon_q/2) + L_{new}(\tau, \varepsilon_q/2) \leq S^q(\tau + d_q) \quad , \text{ for } 0 \leq \tau \leq B_{max} \quad (9)$$

[0063] In this section we show our CAC, which depends on the measurement-based envelopes ($L_N(\cdot; \varepsilon)$ and $S(\cdot)$). In the next section we show how to obtain these envelopes.

Probabilistic Guarantees with Global Effective Envelopes and Service Curves

Deleted: envelopes

[0064] In this section, we derive the probability of delay violation based on the upper bound of the delay shown in the previous section. We then obtain an admission control condition for a single service class by applying the definition of global effective envelope presented in Boorstyn, et al. We finally discuss the admission control conditions for multiple service classes.

[0065] Again we will first consider a single service class and drop the index 'q'. An aggregate of arriving traffic from the same service class in the interval

Deleted: we

(t_1, t_2) , $A_N(t_1, t_2)$ has the same QoS requirements: probability of delay violation ε and delay bound requirement d . $A_N(t_1, t_2)$ is characterized by a family of nonnegative random processes, which have the following properties: (1)

Deleted: has

Stationarity: the statistical properties of $A_N(t_1, t_2)$ do not change with time, and (2)

Independence: The $A_N(t_1, t_2)$ are stochastically independent among all service

Deleted: again

classes. Consider an interval of length of the maximum busy period β . Let $D_N(\tau)$

again denote the ingress-to-egress delay experienced by a bit of an aggregate traffic of the N flows arriving at the domain at time t during the busy period. We assume that $D_N(t)$ is a random variable.

Deleted: which

[0066] For a given delay bound d , the domain guarantees that the ingress-to-egress delay of any arriving bit will not exceed the delay bound d [1.1]; that is, from Eqn. (2.8), for all τ , $0 \leq \tau \leq \beta$,

$$D_N(t) = \min\{z : z \geq 0 \text{ and } \sup_{0 \leq \tau \leq \beta} A_N(t, t+\tau) \leq S(\tau+z)\} \leq d \quad (3.1)$$

or, for all τ , $0 \leq \tau \leq \beta$,

$$\sup_{0 \leq \tau \leq \beta} A_N(t, t+\tau) = E_N(\tau, \beta) \leq S(\tau+d). \quad (3.2)$$

The delay violation occurs if $\exists D_N(t)$ such that

$$D_N(t) \geq d$$

or, $\exists t, \tau$ such that

$$\sup_{0 \leq \tau \leq \beta} A_N(t, t+\tau) > S(\tau+d)$$

With Eqn. (3.2), the probability that the arriving traffic experiences the delay violation is required to be bounded by ϵ and we then have

$$\Pr[D_N(t) \leq d, \forall t, \tau] = \Pr[E_N(\tau, \beta) \leq S(\tau+d), \text{ for all } \tau, 0 \leq \tau \leq \beta] \leq 1 - \epsilon \quad (3.4)$$

[0067] [A] Suppose a set of new flows requests the same service guarantees from a DS domain. The CAC uses the admission control test to ensure that the requested service guarantees and other flows from all service classes can be simultaneously satisfied. From Eqn. (3.4) [1.1] we can obtain the admission

Deleted: [0067]

Deleted: request

Formatted: Font: Not Bold

control condition for a single service class by using the definition of global effective envelope as follows:

Definition of Global Effective Envelope

[0068] A global effective envelope for an interval of length β is a subadditive function $G_N(\tau, \beta, \epsilon)$ that satisfies:

$$\Pr[E_N(\tau, \beta) \leq G_N(\tau, \beta, \epsilon), \forall 0 \leq \tau \leq \beta] \geq 1 - \epsilon$$

It can be seen that a global effective envelope provides a probabilistic bound for the aggregate traffic $A_N(t, t + \tau)$ for every time interval of length τ in β .

Lemma

[0069] A set of flows is partitioned into M classes, and each class contains N_m flows with aggregate arrivals A_{N_m} . Let $G_{N_m}(\tau, \beta, \epsilon)$ be the global effective envelope for class m . Then the following inequality holds: given x ,

$$\text{If } \sum_m G_{N_m}(\tau, \beta, \epsilon) \leq x, \text{ then for all } t, \Pr\left[\exists m \exists \tau: \sum_m E_{N_m}(t, t + \tau) > x(\tau)\right] \leq M \cdot \epsilon.$$

We divide flows from the same service class into two categories, i.e., $M = 2$. The first group is the set of flows, N_1 , that are already admitted to the domain. The second group is the new set of flows, K , requesting the QoS from the DS domain. We denote G_N as the global effective envelope of the first group and G_{new} as the global effective envelope of the second group.

Admission Control Condition for a Single Service Class

[0070] Using the definition of the global effective envelope and the lemma, from Eqn. (3.4)[[.]] we have that the arriving traffic has a delay violation with probability $< \epsilon$ if

$$G_N(\tau, \beta, \epsilon/2) + G_{new}(\tau, \beta, \epsilon/2) \leq S(\tau + d) \quad , \text{ for } 0 \leq \tau \leq B \quad (3.5)$$

[0071] If, for the same service class, the set of new flows are regulated by different deterministic envelopes A_j^* , each set of flows sharing the same A_j^* require a different envelope $G_{new}^j(\cdot; \beta, \epsilon_j)$, where $\sum_j \epsilon_j = \epsilon/2$. Therefore Eqn. (3.5) becomes:

$$G_N(\tau, \beta, \epsilon/2) + \sum_j G_{new}^j(\tau, \beta, \epsilon_j) \leq S(\tau + d) \quad , \text{ for } 0 \leq \tau \leq B \quad (3.6)$$

[0072] The CAC in Eqn. (3.5) does not depend on the topology of the domain and the scheduling algorithms within the domain, but the measurement-based traffic-characterized envelopes, which can be measured at an ingress and an egress node in a DS domain.

Constructing the Global Effective Envelope

[0073] We consider a set of admitted flows, N . We characterize the distribution of the aggregate traffic using the Central Limit theorem, which does not require the knowledge of the underlying distributions of each flow. An approximate global effective envelope $G_N(\tau, \beta, \epsilon)$ of the aggregate traffic can be obtained from the Central Limit Theorem as follows. From the definition of the global effective

envelope, we set $\Pr[E_N(\tau; \beta) \geq x, \forall 0 \leq \tau \leq \beta] \approx \varepsilon$. For given τ_0, β we then have, for every integer k ,

$$G_N(\tau; \beta, \varepsilon) \approx \alpha \overline{X((k+1)\tau/k)} + z' \sigma((k+1)\tau/k), \quad \forall \tau \in [\tau_0, \beta] \quad (3.7)$$

where $\alpha = 1 + 1/(k+1)$, z' has the approximate value $z' = \sqrt{|\log(2\pi\varepsilon')|}$ with $\varepsilon' = \tau_0(\alpha-1)\varepsilon/\beta k$, and $\overline{X(\tau)}$ and $\sigma^2(\tau)$ are measured average and variance of the amount of aggregate traffic, respectively. We will discuss those measurements in the next section.

[0074] From Eqn. (3.4) it can be seen that there are several possible global effective envelopes. For given τ_0, β , the method recursively calculates the Deleted: calculate k_i, α_i , and τ_i for $1 \leq i \leq n$, where τ_0 is the first point which is greater than β [1.1]. We Deleted: therefore w can therefore obtain n points of a global effective envelope as follows:

$$k_i = z \left(z + \frac{\overline{X(\tau_{i-1})}}{\sigma(\tau_{i-1})} \right), \quad (3.8)$$

$$\alpha_i = 1 + \frac{1}{k_i + 1}, \quad (3.9)$$

$$\tau_i = \alpha_i \tau_{i-1}, \quad (3.10)$$

and, from Eqn. (3.7),

$$G_N(\tau_i; \beta, \varepsilon) \approx \alpha_i [\overline{X((k_i+1)\tau_i/k_i)} + z' \sigma((k_i+1)\tau_i/k_i)] \quad (3.11)$$

where $z \approx \sqrt{|\log(2\pi\varepsilon)|}$.

[0075] For a set of new flows, one can obtain the global effective envelope Deleted: [0075] of the new flows G_{new} by using the same recursive method as follows. However,

instead of using the measurement-based $\overline{X(\tau)}$ and $\sigma^2(\tau)$ in Eqs. (3.8) and (3.11), $\overline{X(\tau)} = K\rho_K\tau$ and $\sigma^2(\tau) = K\rho_K\tau(A_K^*(\tau) - \rho_K\tau)$, where ρ_K is the long-term average and $A_K^*(\tau)$ is a deterministic envelope for each new flow.

Admission Control Condition for Multiple Service Classes

[0076] The admission control condition in Eqn. (3.2) does not guarantee that QoS requirements of other service classes using the same ingress-to-egress path will be satisfied if a new set of flows is admitted into the domain. To be able to provide their guarantees, without loss of generality, we assume that the new flows have the highest priority and we test the admission control condition for all the service classes. This can be shown as follows. For each service class we obtain an arriving aggregate traffic envelope $G_N^q(\tau, \beta, \epsilon_q/2)$ and a service envelope $S^q(\cdot)$.

From Eqn. (3.5), new flows, which have an arriving envelope $G_{new}(\tau, \beta, \epsilon_q/2)$, will be admitted if the following condition is satisfied: for all service classes q ,

$$G_N^q(\tau, \beta, \epsilon_q/2) + G_{new}(\tau, \beta, \epsilon_q/2) \leq S^q(\tau + d_q), \text{ for } 0 \leq \tau \leq B \quad (3.12)$$

In this section we show our CAC, which depends on the measurement-based envelopes ($G_N(\cdot; \beta, \epsilon)$ and $S(\cdot)$). In the next section we show how to obtain these envelopes.

Measurements of Perspective Envelopes

[0077] In this section, we show how to measure the characteristic parameters for constructing the global effective envelopes of the arriving aggregate

traffic and the service envelope of the departing aggregate traffic for a single service class. First we measure the average and variance of the amount of arriving aggregate traffic, which are used in Eqn. (3.3). We then measure the service envelope of a single service class traffic. At the end of this section we show the admission control algorithm that we will use in the evaluation example below.

Measurement of Aggregate Arriving Envelopes

[0078] We are interested in the aggregate characteristic of an arriving traffic at ingress node. To characterize the aggregate traffic as a Gaussian process, we need to measure the first and second moment of the amount of traffic arriving at the ingress node. Let $A_N(t_1, t_2)$ denote the arriving traffic from an aggregate flow in the interval (t_1, t_2) . Consider time to be slotted with equal length τ [I.I]; i.e., for MPEG sources, we can think of τ as the frame time. We also divide time into M large intervals whose length is T (see Figure 5). For the m^{th} interval we obtain the bound of an arriving traffic [I.I] X_k^m as a function of time interval $k\tau$. The bounded amount of the measured traffic over the m^{th} interval of length T from the current time t can be obtained as follows:

$$X_k^m = \max_{0 \leq n \leq \left(\frac{T}{\tau} - k\right)} A[t - mT + n\tau, t - mT + (n+k)\tau] \quad (4.1)$$

for $m = 1, \dots, M$ and $k = 1, \dots, \frac{T}{\tau}$.

Deleted: We note that in, the average and variance of arriving traffic rate are measured, while here we measure the average and the variance of the amount of an arriving traffic. Again let

Deleted: the

Deleted: We consider at the current time t .

Deleted: Then the

[0079] Thus for every T time slot, the bounded envelope X_k^m is obtained.

Deleted: [0079]

From these envelopes we can obtain the average \bar{X}_k and the variance σ_k^2 of the measured envelopes as follows:

$$\bar{X}_k = \sum_{m=1}^M X_k^m \quad (4.2)$$

and

$$\sigma_k^2 = \frac{1}{M-1} \sum_{m=1}^M (X_k^m - \bar{X}_k)^2 \quad (4.3)$$

[0080] τ , T , and M are the key parameters for the measurement. An improper choice of τ , T , and M may lead to inaccurately characterizing the arriving aggregate traffic and consequently overestimating the bandwidth required by connections. We varied these parameters in our approach and found that our approach is not very sensitive to such changes. Further discussion of how to choose the parameters can be found in C. Cetinkaya and E. Knightly[1], Egress Admission Control[1], In *Proceedings of IEEE Info. Comm. 2000*, Tel Aviv, Israel, March 2000; and D. Tse and M. Grossglauser[2], Measurement-Based Call Admission Control: Analysis and Simulation[2], In *Proceedings of IEEE Info. Comm. 1997*, Kobe, Japan, April 1997.

Deleted: [0080]

Deleted: Remark

Deleted: Choosing a

Deleted: e

Deleted: e

Deleted: the changing

Deleted: The I

Deleted: INFO

Deleted: COMM

Deleted: INFOCOM

Deleted: Figure

Deleted: 6

Deleted: (

Deleted: a

Deleted:)

Deleted: (

Deleted: b

Deleted:)

Deleted: 6

Deleted: (

Deleted: a

Deleted:)

Deleted: .

Deleted: W

Deleted: file

Deleted: (

Deleted:)

[0081] We provide an example of an arrival envelope from a simulation experiment in Figures 6a and 6b. In Figure 6a, we plot the amount of arrival aggregate traffic as a function of time interval, while Figure 6b depicts the arrival aggregate rate as a function of time interval. The simulation consists of 120 multiplexed independent MPEG-2 video traces. These traces are obtained from the

movie "Starship Troopers" ("Starship"). *Starship* has a peak rate of 18.6336 Mbps and an average rate of 4.26 Mbps. The frame rate of *Starship* is 30 frames per second. This type of traffic has a large peak-to-mean ratio, indicating significant burstiness. From Figure 6, the arriving envelope and arrival rate are significantly smaller than the peak envelope and peak rate, respectively. This is due to the statistical multiplexing of the traffic aggregate. As the figures show, given a long time interval, the amount of the arriving traffic and its arrival rate decrease toward the average rate.

Deleted: the

Deleted: the

Deleted: quite

Deleted: which indicates its

Measurement of Aggregate Service Envelopes

[0082] Consider traffic arriving at a DS domain at an ingress node and leaving at an egress node. For measuring service envelopes, we assume that the clocks in the domain are synchronized. We consider a packet system where packets are serviced at discrete times rather than continuous times. Each packet arriving at the ingress node will be time-stamped. Hence, at the egress node, we are able to determine the ingress-to-egress delay of each packet. Let a_j denote the arriving time and d_j denote the departure time of the j^{th} packet. We consider this traffic to be backlogged whenever at least one packet remains in the system. The backlogging condition can be checked by comparing the arriving and the departure times of packets. Traffic is continuously backlogged for k packets in the interval $[a_j, d_{j+k-1}]$ if

Deleted: at the

Deleted:

$$d_{j+m} > a_{j+m+1}, \text{ for all } 0 \leq m \leq k-2 \quad (4.4)$$

for $k \geq 1$.

[0083] Next, let *ServEnv* be a list containing $(x, S_j^k(x))$ pairs which represent the amount of time $S_j^k(x)$ required to service x bits when considering the packet j with k backlogged packets. Again consider an interval of length T , which contains the total number of arriving packets n . For packet j , we consider the delay of the packet and also the delay of the packets backlogged after this packet. We thus can define $S_j^k(x)$ as follows:

$$S_j^k(x) = d_{j+k-1} - a_j \quad (4.5)$$

for all $k \geq 1$ which satisfy the backlogging condition.

A bounded service envelope can be obtained as follows:

```

Initialize (ServEnv)
for packet  $j = 1$  to  $n$ 
     $S_j^1 \leftarrow d_j - a_j$ 
     $x(S_j^1) \leftarrow \text{Size}(\text{packet}(j))$ 
     $k \leftarrow 2$ 
    while ( $d_{j+k-1} > a_{j+k-1}$ )
         $S_j^k \leftarrow d_{j+k-1} - a_j$ 
         $x(S_j^k) \leftarrow x(S_j^{k-1}) + \text{Size}(\text{packet}(j+k-1))$ 
         $k \leftarrow k + 1$ 
    Merge (ServEnv,  $(x(S_j^k), S_j^k) \forall k \in \text{ServEnv}$ )
for  $i = (\text{Length}(\text{ServEnv}) - 1)$  to 0
    if ( $x[i] < x[i-1]$ )
        Remove ( $x[i-1], S[i-1]$ )
Using ServEnv to plot a service envelope.

```

Table I Bounded Service Algorithm

[0084] Figure 7a shows a service envelope of an aggregate of departing traffic obtained from a simulation. The simulation consists of a node with a 622 Mbps link and the same number of sources (*Starship*) used for Figure 6a and 6b. It

is noted that the envelope is an increasing function. In addition, Figure 7b depicts the slope of the service envelope. The slope of the envelope indicates variation of the service rate for the burstiness of the arriving aggregate flows.

Deleted: 7

Deleted: (

Deleted: b

Deleted:)

Admission Control Algorithm using local effective envelopes

- I. To request the QoS for its traffic from a DS domain, the user submits to a BB its traffic characterization containing $A_j^*(x)$, QoS requirements, d and ϵ .
- II. CAC constructs a local effective envelope, L_{new} , based on the submitted parameters by using Eqs. (7) and (8).
- III. CAC then uses the measurement of the average and variance of arriving traffic to construct local effective envelopes for admitted flows, $L_{N_s}^q$, from Eqn. (4).
- IV. CAC finally makes an admission control decision based on L_{new} , $L_{N_s}^q$, and S'' obtained from the previous section. The request is accepted if the condition in Eqn. (9) is true.

Deleted: another

Deleted:

Deleted: Section as follows

Admission Control Algorithm using Global Effective Envelope

- I. To request the QoS for its traffic from a DS domain, the user submits to a BB its traffic characterization containing $A_j^*(x)$, QoS requirements, d and ϵ .

- ii. CAC constructs a global effective envelope, G_{new} , based on the submitted parameters by using Eqs. (3.8) – (3.11).
- iii. CAC then uses the measurement of the average and variance of arriving traffic (obtained from Eqs. (4.2) and (4.3) respectively) to construct global effective envelopes for admitted flows, $G_{N_i}^g$, from Eqs. (3.8)-(3.11).
- iv. CAC finally makes an admission control decision based on G_{new} , $G_{N_i}^g$, and S^g obtained from the previous section. The request is accepted if the condition in Eqn. (3.12) is true.

Deleted: another

Deleted:

Deleted: Section as follows

Evaluation and Simulation Results

[0085] In this section, we evaluate the effectiveness of our CAC condition.

We consider a case where all flows are homogeneous, i.e., all flows satisfy the same deterministic envelope and require the same delay QoS guarantee from a DS domain. *Starship* is again used as a traffic source. For simulations, we set the probability of delay violation $\varepsilon = 10^{-6}$ and vary the delay bound requirement d . We consider a DS domain, which has an egress node with a link capacity of 622 Mbps ($C = 622$ Mbps for all simulations). All links in the domain have the same capacity as the egress link. Each node has a FIFO scheduler. We consider three scenarios where the number of nodes and amount of cross traffic in the domain are changed as follows:

Deleted:

Deleted: that

Deleted: that is

Deleted: has

Deleted: of

Deleted: is

Deleted: 8

Deleted: (

Deleted: a

Deleted:)

Deleted: 8

Deleted: (

Deleted: b

Deleted:)

- I. The DS domain consists of a node without any cross traffic (see Figure 8a).
- II. The DS domain consists of two nodes without any cross traffic (see Figure 8b).

III. The DS domain consists of three nodes with some cross traffic (see Figure 8c).

Note that in configuration III, we use 40 flows of *Starship* as cross traffic requesting the same QoS requirements as the main traffic.

[0086] We will evaluate our approach ('Approach') by using the CAC algorithm described above. As benchmarks for our approach, we also investigate two of the following approaches and a set of simulations ('Simulation'). First, we consider average rate allocation ('Average Rate'), where the maximum number of admissible flows is determined by the link capacity divided by the long-term average rate of each flow and, therefore, the average link utilization for this approach is unity. Second, we consider the peak rate allocation ('Peak Rate'), where the maximum number of admissible flows is determined by the link capacity divided by the peak rate of each flow. For simulation purposes, each flow starts randomly sending traffic into the DS domain. Many simulations are performed. The average result of the simulation is presented. We compare the average link utilization, U_{ave} , obtained from each approach. U_{ave} is defined as follows:

$$U_{ave} = \frac{\sum_{i=1}^N \rho_i}{C}$$

where N is the number of connections admitted to the domain, ρ_i is the long-term average rate of each flow, and C is the link capacity at the egress node.

[0087] Figures 8a-8c show the three scenarios we investigate and the corresponding results where the average link utilization is a function of the delay

Deleted: g
Deleted: (
Deleted: c
Deleted:)
Deleted: also
Deleted:
Deleted: with
Deleted: of
Deleted: the

Deleted: the
Deleted: proposes

Deleted: connection

Deleted:
Deleted: as

bound. From the results, the link utilization obtained by the simulations and our approach is considerably higher than the one obtained by the peak rate allocation.

This shows that the simulations and our approach can exploit the multiplexing gain from the traffic aggregate and, hence, yield high link utilization. For example, from Figure 8a, for the delay bound $d = 10$ ms, the simulation and our approach yield 86% of the link utilization while the peak rate allocation can only achieve 82%. In addition, the results show that our approach can accurately characterize arriving and departing aggregate traffic. Figures 8b-8c show that our approach can determine an accurate condition of the domain without the knowledge of the underlying topology and cross traffic, and yield link utilization close to simulation results.

[0088] From the foregoing it will be appreciated that the measurement-based admission control algorithm exploits global effective envelopes and service envelopes to accurately characterize arriving and departing aggregate traffic. The CAC algorithm is tunable through a set of parameters. The algorithm is scalable and can support a variety of service classes. We showed the effectiveness of the CAC algorithm by comparing the link utilization obtained by the algorithm to results from simulations. While the invention has been described in its presently preferred form, it will be understood that the principles of the invention may be applied to many applications and that implementation details of the algorithm can be varied without departing from the spirit of the invention as set forth in the claims.

Deleted: is depicted

Deleted: the

Deleted: Fig 8

Deleted: (

Deleted: a

Deleted:)

Deleted: ec

Deleted: , respectively, W

Deleted: obtain

Deleted:

Deleted: 8

Deleted: (

Deleted: b

Deleted:)

Deleted: -

Deleted: (

Deleted: c

Deleted:)

Deleted: obtain the

Deleted: the

Deleted: the

Deleted: present invention provides a

Deleted: that

Deleted: the

Deleted: aggregate

Deleted: Based on the measurement, t

Deleted: the

Deleted: the

Deleted: may

Deleted: application

Deleted: appended